

GGF-Grid Interoperability Now (GIN) SRM Interoperability Testing Results

May 10, 2006

Alex Sim, Vijaya Natarajan
Scientific Data Management Research Group
Lawrence Berkeley National Laboratory

Abstract:

Storage Resource Managers (SRMs) are based on a common interface specification (<http://sdm.lbl.gov/srm-wg>). SRMs can have different implementations and services for the underlying storage systems. The purpose of the GIN SRM testing is to check the compatibility and interoperability of SRMs according to the specification.

The tests include the following:

- read access to a file in a remote data grid managed by SRM
- write access to a registered user account in a remote data grid managed by SRM
- a file replication for a registered user between two independent data grids
- space reservation and write access to the reserved space for a registered user in a remote data grid by managed SRM (for SRM v2.1)

Approach:

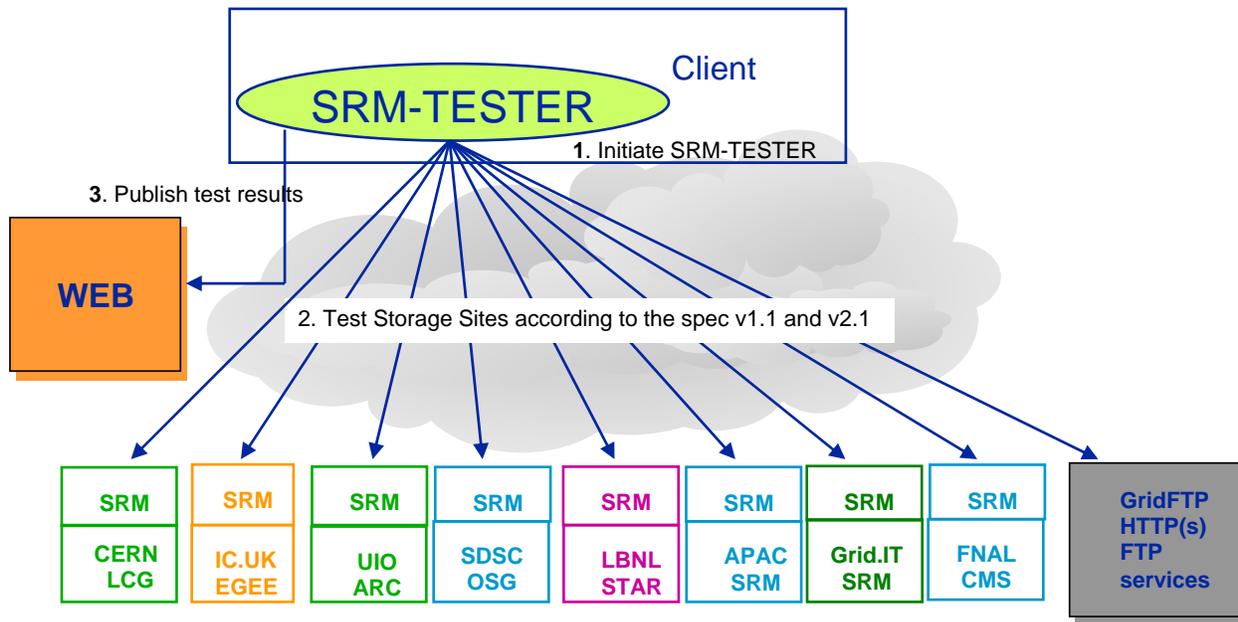
Participating storage sites that support SRM access will be tested with the SRM-Tester for various operations, according to the SRM specification. In case of failures, other client tools may be used to test the SRM servers further.

SRM-Tester:

The SRM-Tester tests Storage Sites managed by SRMs. It is an independent implementation of SRM specification for the purpose of testing. It is designed to:

- test the conformity of the SRM server interface according to the SRM spec v1.1 and v2.1
- test the compatibility and interoperability of the SRM servers according to the SRM specifications
- test the following operations for SRM v1.1:
 - ping, put, get, advisory delete, copy
- test the following operations For SRM v2.1:
 - ping, put, get, release file, copy, reserve space, release space, remove
- test the following file transfer protocols:
 - gsiftp, ftp, http, https
- publish the testing results to a web site.

The test results for GIN-Data SRM testing are published at <http://sdm.lbl.gov/srm-tester/ggf17.html>.



Participating Sites:

There are 7 sites that participated in the GIN-data SRM island testing. 5 SRM V1.1 and 2 SRM V2.1 were tested. SRMs at 2 sites were not ready at the time of the testing.

- APAC
 - SRM is not ready due to undergoing SFT tests
- ARC
 - `srm://grid.uio.no:58000`
- EGEE
 - `srm://gfe02.hep.ph.ic.ac.uk:8443`
- Grid.it
 - SRM is not installed publicly and not ready for testing
- LCG/EGEE
 - `srm://lxdpm01.cern.ch:8443`
 - `srm://lxdpm01.cern.ch:8444 (SRM v2.1)`
- OSG
 - `srm://t2data2.sdsc.edu:8443`
- STAR
 - `srm://dmx09.lbl.gov:8443`
 - `srm://dmx09.lbl.gov:8445 (SRM v2.1)`

SRM at FNAL was not participating in this activity, but with help from Timur Perelmutov at FNAL, we were able to use SRM at FNAL as a backup SRM for additional testing with non-VO access.

- FNAL
 - `srm://fledgling05.fnal.gov:8443//srm/managerv1`
 - `srm://fledgling05.fnal.gov:8443//srm/managerv2 (SRM v2.1)`

Information requested for SRMs:

1. SRM end point:
e.g. in case of an SRM at SDSC, it looks like
srm://t2data2.sdsc.edu:8443//srm/managerv1
2. Version of SRM running on the above end point:
e.g. v1.1 or v2.1
3. writable data directory for GIN-VO
e.g. /data2/gin
4. if "delete" on the file in SRM managed cache is available
e.g. true, false
5. if "overwrite" on the existing file in SRM managed cache is available
e.g. true, false
6. if SRM managed cache is "read-only" storage. In other words, no writing is possible.
in case of read-only cache, the URL of the file that we can pull out
e.g. srm://t2data2.sdsc.edu:8443//srm/managerv1?SFN=/data2/gin/mytest.file
7. supported transfer protocol list
e.g. http, gsiftp, dcap

From the above information, we could compose an URL and proceed with tests:
srm://t2data2.sdsc.edu:8443//srm/managerv1?SFN=/data2/gin/my_put_test.file

Testing Description for SRM v1.1

1. "ping" the target SRM server to check if the server is running
 - a. GSI security may or may not be enforced.
2. "put" (write) a file into the SRM managed, VO specific directory from client's local disk
 - a. checks the transfer protocol during the file transfers
3. "get" (read) a file from SRM managed, VO specific directory to the client's local disk
 - a. Requests a file from the SRM managed, VO specific directory
 - b. When the file is ready, SRM server returns Transfer URL (TURL) with transfer protocol
 - c. File transfers from the TURL to the client's local disk is tested
4. "unpin" (release) files
5. "advisory delete" (remove) files
6. "copy" a file from one SRM to another SRM
 - a. SRMs are in two separate independent data grids
7. "copy" a file from one remote gridftp server to an SRM

Testing Description for SRM v2.1

1. "ping" the target SRM by "srmGetRequestIDRequest"
 - a. V2.1 does not have an explicit "ping" interface
2. "put" (write) a file into SRM managed, VO specific directory from client's local disk
 - a. Checks the transfer protocol during the file transfers
 - b. Tests srmPutDone() after file transfer is completed
3. "get" (read) a file from SRM managed, VO specific directory to client's local disk
 - a. Requests a file from the SRM managed, VO specific directory
 - b. When the file is ready, SRM server returns Transfer URL (TURL) with transfer protocol
 - c. File transfers from the TURL to the client's local disk is tested
4. "release" files
 - a. releases a file after file transfer is completed (following "get" operation)
5. "non-recursive directory listing" (ls)

- a. browses a file or a directory in the SRM managed, VO specific space
6. "copy" a file
 - a. Request to "cp" a file from one space to another space in the same SRM
 - b. Request to "replicate" a file from one SRM to another SRM. SRMs are in two separate independent data grids.
 - c. Request to "replicate" a file from a remote gsiftp server to an SRM.
7. "reserve a space"
 - a. requests to reserve a volatile or durable space under the SRM for an amount of byte size and gets a corresponding space token
8. "status" of the reserved space
 - a. checks the status of the space with the space token received from the above space reservation
9. "put" (write) a file into the above reserved space with the space token
 - a. Similar to "put" operation, but a specific space token is used.
10. "release a space"
 - a. releases the above reserved space with the space token

Test results for SRM v1.1

Table 1 : Test results for SRM v1.1

	ping	get	put	advisory delete
EGEE (IC.UK)	pass	pass	pass	pass
LCG/EGEE (CERN)	pass	pass	pass	pass
OSG (SDSC.EDU)	pass	pass	pass	pass
STAR (LBL.GOV)	pass	pass	pass	pass
ARC (UIO.NO)	pass	fail	fail	fail

- SRM at ARC fails operations except "ping" during the status checking
 - Throws exception : java.lang.NumberFormatException
 - SRMCP (v1.23) has the same error
 - NGCP (Nordugrid ARC standalone v0.5.46) works for operations except copy
 - ngcp does not handle different formats of certificates yet, such as service certificates

COPY between SRMs in PULL mode

- SRM-Tester sends a request to the target SRM to "pull" a file from the source SRM
- After a copy request is submitted to the target SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.
- When the first copy test failed, SRM at FNAL was used as the source and target SRM for additional testing, and testing with srmcp client tool was performed.

Table 2 : Test results for COPY between SRMs in PULL mode

Target SRM →	EGEE (IC.UK)	LCG/EGEE (CERN.CH)	OSG (SDSC.EDU)	STAR (LBL.GOV)	ARC (UIO.NO)
Source SRM ↓					
EGEE (IC.UK)		N.A.	N.A.	pass	fail
LCG/EGEE (CERN.CH)	pass		N.A.	pass	fail
OSG (SDSC.EDU)	pass	N.A.		pass	fail
STAR (LBL.GOV)	pass	N.A.	N.A.		fail
ARC (UIO.NO)	fail	N.A.	N.A.	fail	

- SRM at CERN (LCG/EGEE)
 - The request goes through the interface, and the status shows request queued and in progress.
 - We were told that actual copy is not implemented yet
- SRM at SDSC (OSG)
 - The request goes through the interface, and the status shows request queued and in progress, with no further explanation
 - We were told that only “push” mode into the SRM at SDSC works because the underlying storage resides inside the private network.
- SRM at UIO (ARC)
 - The request doesn't go through.

COPY between SRMs in PUSH mode

- SRM-Tester sends a request to the source SRM to “push” a file from the source SRM to the target SRM
- After a copy request is submitted to the source SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.
- When the first copy test failed, SRM at FNAL was used as the source and target SRM for additional testing, and testing with srmcp client tool was performed.

Table 3 : Test results for COPY between SRMs in PUSH mode

Target SRM →	EGEE (IC.UK)	LCG/EGEE (CERN.CH)	OSG (SDSC.EDU)	STAR (LBL.GOV)	ARC (UIO.NO)
Source SRM ↓					
EGEE (IC.UK)		pass	pass	pass	fail
LCG/EGEE (CERN.CH)	N.A.		N.A.	N.A.	N.A.
OSG (SDSC.EDU)	pass	pass		pass	fail
STAR (LBL.GOV)	pass	pass	pass		fail
ARC (UIO.NO)	fail	N.A.	fail	fail	

- SRM at CERN (LCG/EGEE)
 - The request goes through the interface, and the status shows request queued and in progress.
 - We were told that actual copy is not implemented yet
- SRM at UIO (ARC)
 - The request doesn't go through.

COPY from gsiftp server to SRM

- SRM-Tester sends a request to the target SRM to “pull” a file from the source gsiftp server. Gsiftp server is based on GT 4.0.1.
- After a copy request is submitted to the target SRM, communication thereafter is between the SRM and the source gsiftp server.
- SRM-Tester checks the status of the copy request.
- When the first copy test failed, gsiftp server based on GT 2.4.3 at LBNL (STAR) was used as the source for additional testing, for possible certificate format handling issues.

Table 4 : Test results for COPY from GSIFTP server to SRM

Target SRM →	EGEE (IC.UK)	LCG/EGEE (CERN.CH)	OSG (SDSC.EDU)	STAR (LBL.GOV)	ARC (UIO.NO)
Source GSIFTP ↓					
Gsiftp server from GT4.0.1	pass	N.A.	N.A.	pass	fail

- SRM at CERN (LCG/EGEE)
 - The request goes through the interface, and the status shows request queued and in progress.
 - We were told that actual copy is not implemented yet.
- SRM at SDSC (OSG)
 - The status returned failure with no further explanation
 - We were told that only “push” mode into the SRM at SDSC works because the underlying storage resides inside the private network.
- SRM at UIO (ARC)
 - The request doesn’t go through.

Test results for SRM v2.1

Table 5 : Test results for SRM v2.1

	ping	put	putDone	get	releaseFile	browse (LS)
LCG/EGEE (CERN)	pass	pass	pass	pass	pass	pass
STAR (LBL.GOV)	pass	pass	pass	pass	pass	pass

	Reserve Space	Get Status	Put a file in the reserved space	Release Space
LCG/EGEE (CERN)	pass	pass	pass	pass
STAR (LBL.GOV)	pass	pass	pass	pass

COPY between SRMs in PULL mode

- SRM-Tester sends a request to the target SRM to “pull” a file from the source SRM
- After a copy request is submitted to the target SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.

Table 6 : Test results for COPY between SRMs in PULL mode

Target SRM →	LCG/EGEE (CERN.CH)	STAR (LBL.GOV)
Source SRM ↓		
LCG/EGEE (CERN.CH)		pass
STAR (LBL.GOV)	N.A.	

- SRM at CERN (LCG/EGEE)
 - The request goes through the interface, and the status shows request queued and in progress.
 - We were told that actual copy is not implemented yet

COPY between SRMs in PUSH mode

- SRM-Tester sends a request to the source SRM to “push” a file from the source SRM to the target SRM
- After a copy request is submitted to the source SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.

Table 7 : Test results for COPY between SRMs in PUSH mode

Target SRM →	LCG/EGEE (CERN.CH)	STAR (LBL.GOV)
Source SRM ↓		
LCG/EGEE (CERN.CH)		N.A.
STAR (LBL.GOV)	pass	

- SRM at CERN (LCG/EGEE)
 - The request goes through the interface, and the status shows request queued and in progress.
 - We were told that actual copy is not implemented yet

COPY from gsiftp server to SRM

- SRM-Tester sends a request to the target SRM to “pull” a file from the source gsiftp server. Gsiftp server is based on GT 4.0.1.
- After a copy request is submitted to the target SRM, communication thereafter is between the SRM and the source gsiftp server.
- SRM-Tester checks the status of the copy request.
- When the first copy test failed, gsiftp server based on GT 2.4.3 at LBNL (STAR) was used as the source for additional testing, for possible certificate format handling issues.

Table 8 : Test results for COPY from GSIFTP server to SRM

Target SRM →	LCG/EGEE (CERN.CH)	STAR (LBL.GOV)
Source GSIFTP ↓		
Gsiftp server from GT4.0.1	N.A.	pass

- SRM at CERN (LCG/EGEE)
 - The request goes through the interface, and the status shows request queued and in progress.
 - We were told that actual copy is not implemented yet.

Lessons Learned

1. Preparation
 - a. Setting up a VO (gin.ggf.org) access to the host and the VO designated account was error-prone.
 - i. Email communication was needed with system people
 - ii. However, once it’s set up, no problems
 - b. The necessary information to perform testing can be summarized as the following five items.
 - i. SRM end point: e.g. srm://t2data2.sdsc.edu:8443//srm/manager
 - ii. Version of SRM running on the above end point: e.g. 1.1, 2.1.2
 - iii. writable data directory for GIN-VO: e.g. /data/gin
 - iv. Available “copy” operational type: e.g. PULL, PUSH, ALL
 - v. Optional supported transfer protocol list : e.g. gsiftp, http, ftp
2. SRM v1.1 operations

- a. well-understood by most implementations
 - b. Compatible and interoperable between SRMs
 - c. Problem at ARC SRM at uio.no could be a simple data type issue in the status. Further investigation is needed.
3. SRM v2.1 operations
- a. Two participating implementations
 - i. CERN DPM and LBNL SRM
 - ii. Compatible and interoperable between SRMs
 - b. Additional Non-VO access to SRM at FNAL
 - i. Compatible and interoperable between SRMs

Summary

Most of the SRM v1.1 operations were successful. The problem at ARC SRM at uio.no could be a simple data type issue in the status that requires further investigation. Otherwise, SRM servers conform to the interface specification, and the interoperability was successful.

For SRM v2.1 testing, it was successful between the two participating groups.

In the future, additional implementations will be tested as they become available.

References

1. Grid Interoperability Now (GIN) Working Group, <https://forge.gridforum.org/projects/mgi>
2. GIN VOMS, <https://kuiken.nikhef.nl:8443/voms/gin.ggf.org/>
3. SRM Collaboration, <http://sdm.lbl.gov/srm-wg>
4. SRM Specification v1.1, <http://sdm.lbl.gov/srm-wg/doc/srm.v1.0.pdf>
5. SRM Specification v2.1, <http://sdm.lbl.gov/srm-wg/doc/SRM.spec.v2.1.1.html>
6. SRM-Tester, <http://sdm.lbl.gov/srm-dist>
7. SRMCP, <http://srm.fnal.gov>
8. NGCP, <ftp://ftp.nordugrid.org/nordugrid/tags/0.5.46/>
9. GIN SRM island testing results, <http://sdm.lbl.gov/srm-tester/ggf17.html>

Acknowledgement

- APAC
 - Univ. of Melbourne : Glenn Moloney <glenn@physics.unimelb.edu.au>
- ARC
 - Lund Univ. : Oxana Smirnova <oxana.smirnova@hep.lu.se>
- EGEE
 - CERN : Erwin Laure <Erwin.Laure@cern.ch>
 - Imperial College : Mona Aggarwal <m.aggarwal@imperial.ac.uk>
Olivier van der Aa <o.van-der-aa@imperial.ac.uk>
David Colling <d.colling@imperial.ac.uk>
- FNAL
 - Timur Perelmutov <timur@fnal.gov>
- Grid.IT
 - INFN : Riccardo Zappi <riccardo.zappi@cnaf.infn.it>
Luca Magnoni <luca.magnoni@cnaf.infn.it>
- LCG/EGEE
 - CERN : James Casey <James.Casey@cern.ch>
Jean-Philippe Baud <Jean-Philippe.Baud@cern.ch>
- OSG
 - UCSD : Frank Wuerthwein <fkw@ucsd.edu>
Abhishek Singh Rana <rana@ucsd.edu>
- STAR

- BNL: Jerome Lauret <jeromel@bnl.gov>
 - LBNL : Alex Sim <asim@lbl.gov>
 - Vijaya Natarajan <vnatarajan@lbl.gov>
- GIN VOMS VO
 - NIKHEF : Oscar Koeroo <okoeroo@nikhef.nl>
 - INFN : Vincenzo Ciaschini <vincenzo.ciaschini@cnaf.infn.it>